Pergamon

Stud. Hist. Phil. Sci., Vol. 32, No. 1, pp. 31–67, 2001
© 2001 Elsevier Science Ltd. All rights reserved.
Printed in Great Britain
0039-3681/01 $ - see front matter

www.elsevier.com/locate/shpsa

SETI: On the Prospects and Pursuitworthiness of the Search for Extraterrestrial Intelligence

*André Kukla**

My topic is extraterrestrial intelligence. Following current conventions, I use the abbreviation 'ETI' to stand for three related concepts: (1) the abstract idea of extraterrestrial intelligence, (2) individuals who are both extraterrestrial and intelligent (as in 'There's an ETI in the closet'), and (3) the hypothesis that there are ETIs. SETI is the search for ETIs, and CETI is the attempt to communicate with ETIs. In this paper, I will try to answer the two most basic questions in extraterrestrial studies. First, what is the status of the ETI hypothesis? In the light of what we know, how likely is it that there are ETIs? Secondly, what is the value of pursuing SETI? Are the expected epistemic and pragmatic payoffs sufficient to warrant an extensive search? The first, substantive question is discussed in Section 2, Section 3 and Section 4. The second, 'program evaluation' question is the subject of Section 5 and Section 6. But first we'll take a brief look at some of the early history of the topic. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Extraterrestrials; SETI; Exobiology; Philosophy of Science.

1. A Brief History of ETI

Humans have been concerned with extraterrestrial questions more or less continuously since antiquity. Those who aren't familiar with the history of the subject may be surprised to learn that for the most part, it was the pro-ETI view that there *are* intelligent extraterrestrials that was in ascendancy. There were some dissenting voices—among them St Augustine, St Thomas Aquinas and Albertus Magnus.[1] But they were greatly outnumbered by the pro-ETI forces, which included Aristotle,

* Division of Life Sciences, University of Toronto at Scarborough, Scarborough, Ontario M1C 1A4, Canada (*e-mail*: andre.kukla@utoronto.ca)

Received 11 October 1999; in revised form 28 February 2000.

[1]My historical information is drawn entirely from Beck (1985), Crowe (1986) and Dick (1982). These monographs should be consulted for references to original sources.

PII: S0039-3681(00)00025-X

Plutarch, Lucretius, Nicholas of Cusa, Giordano Bruno, Gassendi, Locke, Lambert, Kant, Thomas Wilkins, Christiaan Huygens and Bernard de Fontenelle.[2]

The two ancient writers who dealt most extensively with the topic of ETI were Plutarch and Lucretius. Both of them were pro-ETI, but each gave different reasons for his position. Their arguments presaged the two basic strategies for arguing in favor of ETI for centuries to come. Some features of their arguments are parts of the pro-ETI armamentarium to this day. Plutarch inferred that intelligent life exists on the moon from four premises: (1) the earth has no privileged position in the universe; (2) the arrangement of the heavens indicates intelligent design; (3) the moon is sufficiently like the earth to support life; (4) if the moon didn't have life, it would exist to no purpose, and this conflicts with the assumption of intelligent design. Of course, we know now that the conclusion of this argument is false. There's no life on the moon—we've been there. But the same teleological argument was made by Bruno and Nicholas of Cusa for the billions of stars: would an intelligent designer have wasted so much space?

This is not the place to argue about the legitimacy of teleological explanations. Grant that the universe has been intelligently designed. Doesn't the empirical refutation of Plutarch's original lunar hypothesis show that the general argument-form is invalid, even when it's applied to galaxies? Apparently the Intelligent Designer *is* willing to build entire worlds without life. There must be something about Her plan that we don't understand. This brings up the question of Plutarch's third premise—that the moon is sufficiently like the earth to support life. The occurrence of this premise in the argument shows that Plutarch didn't suppose that an intelligently designed universe would not contain uninhabited planets. If that were his view, then the third premise would be superfluous. Evidently, what Plutarch's teleological intuition told him was that an intelligent designer wouldn't make *inhabitable* planets and then leave them uninhabited. Let's grant this hypothesis about intelligent design as well. Then the empirically established fact that the moon contains no life doesn't by itself invalidate the later galactic version of the argument—it just shows the premise about the moon to be false. But also, it shows that the bare existence of billions of stars doesn't by itself make the teleological argument work. It's not enough to establish that there's sufficient *space* for other life forms. The Apollo missions to the moon already tell us that the (presumed) intelligent design of the universe must be compatible with the existence of large uninhabited bodies. The Plutarchian argument needs the premise that among the billions of stars, there are other *inhabitable* places. There was nothing about Cusa's or Bruno's bare observation of the teeming stars in the night sky that could have provided them with that premise. Contemporary proponents of SETI think that they have good reasons for supposing that there are many other inhabitable places in the universe. To be sure, they don't embed this hypothesis in a teleological argument. But both

[2]William Whewell belonged successively to both camps. See Crowe (1986).

the Plutarchian argument and the contemporary SETI arguments for the existence of ETIs rely on the hypothesis that there are other inhabitable places. We'll examine the contemporary arguments in due course. If the case for their inhabitability hypothesis should prove to be compelling, it would lend a measure of credence to the Plutarchian argument as well. But the galacticized version of Plutarch's argument by itself is inadequate, even if we accept its teleological assumptions. Maybe the Intelligent Designer made the stars for purely decorative reasons.

There's also a more parochial class of pro-ETI argument from deistic premises. The hypothesis of a plurality of worlds inhabited by intelligent beings poses no special problems for devout Hindus or Buddhists. But it's a central idea of the Judeo-Christian tradition that there's a unique relationship between God and human beings. The Judeo-Christian God spends an enormous amount of time concerning Himself [sic] with human affairs. This observation doesn't by itself entail that Judeo-Christian orthodoxy is incompatible with ETI—for it's possible that God is simultaneously carrying out similar activities throughout the universe—that He processes the needs of His myriads of cosmic children in parallel mode. But it's of the essence of Christianity that God sacrificed His *only* begotten son to save humanity. If the only begotten son of God merely saved *some* of the intelligent beings in the universe, what has been the disposition of the rest? As Thomas Paine noted, the existence of a plurality of worlds renders Christianity 'little and ridiculous, and scatters it in the mind like feathers in the air' (Paine, 1948, p. 44). What follows from this observation? As the saying goes, one person's *reductio ad absurdum* is another's discovery of a novel consequence. That a plurality of worlds renders Christianity little and ridiculous can be regarded as an argument against the plurality of worlds, or as a proof of the smallness and silliness of Christianity.

Lucretius' argument is much more modern than Plutarch's. For one thing, he dispenses entirely with teleological assumptions. According to Lucretius, the conclusion of ETI follows from purely mechanical assumptions, together with the premise that the universe is infinite in extent:

> Granted, then, that empty space extends without limit in every direction and that seeds innumerable in number are rushing on countless courses through an unfathomable universe under the impulse of perpetual motion, it is in the highest degree unlikely that this earth and sky is the only one to have been created . . . This follows from the fact that our world has been made by nature through the spontaneous and casual collision and the multifarious, accidental, random and purposeless congregation and coalescence of atoms whose suddenly formed combinations could serve on each occasion as the starting-point of substantial fabrics—earth and sea and sky and the races of living creatures. (Lucretius, 1951, pp. 91–92)

What replaces design here is the same thing that Darwin replaced it with in his account of the origin of species: chance. The argument is that if the world is an infinite collection of atoms jostling about at random, then any conglomeration whatever is going to occur many times over. The same argument-form is employed in drawing the conclusion that a roomful of monkeys hitting typewriter keys at ran-

dom will eventually produce the verbatim text of *Hamlet*—and that if we wait long enough, they'll produce it again. This type of mechanical argument is not merely still around—somewhat amended versions of it are presently the mainstay of the pro-SETI case. We'll see that in some ways, Lucretius' original argument is more robust than some of its descendants. The vicissitudes of this argument are part of the modern, scientific story, and will be dealt with at a later time. In any case, whether from teleological or mechanistic assumptions, the existence of ETIs was the received view by 17th and 18th centuries, and has remained so, though there have been notable dissenting voices.

Plutarch's and Lucretius' arguments bear on the substantive question whether there are any ETIs. History also records the voices of some who objected, not to the ETI hypothesis, but to the idea that the ETI hypothesis raises an interesting or important question. Montaigne, Milton and Pope held the view that speculation abut ETIs is a waste of time and that we would do better to think about more mundane issues. Let's call the question whether the ETI hypothesis is worth pursuing the *metaquestion* of extraterrestrial studies. I'll deal with both the substantive question and the metaquestion in the next few sections, starting with the substantive issue. This may seem like putting the cart before the horse—for shouldn't we ascertain that the substantive question is worth an answer before we try to answer it? It would be nice always to know whether a project is pursuitworthy before we invest an iota of time or energy in its pursuit. But this ideal is unrealizable. In fact, it leads to an infinite regress—for wouldn't we have to know that the *meta*-question (i.e., whether the ETI hypothesis is worth pursuing) is itself pursuitworthy before devoting an iota of effort to it? In practice, there's no way to avoid an initial speculative investment of effort. So we might just as well speculate that the ETI question is pursuitworthy as that the metaquestion is pursuitworthy.[3]

Moreover, the resolutions of both questions are inextricably intertwined. If the discovery of ETIs promises to bring enormous benefits in its train, that would be a reason to devote large quantities of resources to looking for them. But this reason would be negated if the probability of there *being* any ETIs to contact were vanishingly small. So the answer to the metaquestion depends in part on the answer to the substantive question. The most reasonable way to proceed is, I think, clear and uncontentious. Before devoting millions more to the search for ETIs, we need a low-cost, paper-and-pencil initial assessment of the prospects for the success of such a project. It's true that to do even this much is to presume that the pursuit of extraterrestrial issues has some value. But the working time of scholars is cheaply bought. Surely everyone will agree that the pursuit of the ETI hypothesis is worth at least a few weeks of a few scholars' time, plus stationery—say twenty thousand 1999 US dollars. If the result of this cheap initial assessment of the ETI hypothesis is that there's no chance of our ever finding out that there are ETIs,

[3]A more elaborate version of the foregoing argument can be found in Kukla (1991).

then that may be the end of the story. If the result of the initial assessment is anything else, then it becomes one factor among others in the determination of whether larger-scale investigations of the ETI question are a good investment.

2. Inductive Arguments for ETI

How plausible is it to suppose that there are ETIs? Discussions of this issue are generally guided by the frequently-cited Drake equation.[4] There are several minor variants of this equation in the ETI literature. Mash gives the following version:

$$N = N_* f_p n_e f_l f_i f_e f_L,$$

where

> N = the number of presently existing advanced civilizations in the galaxy (where 'advanced' means 'at least capable of, and engaged in interstellar communication');
> N_* = the number of stars in the galaxy;
> f_p = the fractions of those stars with planets;
> n_e = the average number of planets in the solar system whose environments are suitable for life;
> f_l = the fraction of habitable planets on which life occurs;
> f_i = the fraction of life-bearing planets on which intelligence evolves;
> f_c = the fraction of intelligence-bearing planets on which advanced civilizations arise; and
> f_L = the fraction of advanced civilizations existing at this time (as determined by the average lifetime of a civilization divided by the present age of the galaxy, or, more precisely, the age of the civilization-bearing galaxy). (Mash, 1993, pp. 210–211)

The last two factors, f_c and f_L, are interestingly different from the rest. As is noted in the definition of N, 'advanced' civilizations are those that have the technological capability for sending and receiving interstellar signals. Of course, everyone concedes that it's possible to be intelligent without possessing this capability. We didn't possess it ourselves until quite recently. So why the restriction? The reason, as Mash notes, is that the Drake equation is

> oriented toward the funding of SETI projects, that is, projects that intend to *search* for extraterrestrial civilizations, rather than toward the more academic question of whether other civilizations might exist anywhere at any time. (p. 211)

So we need to distinguish two substantive hypotheses: the *ETI hypothesis* that there *are* ETIs, and the *SETI hypothesis* that we can *make contact with* ETIs. It's conceivable that a general argument based on firmly established astrophysical, biological and psychosocial principles might convince us that it's overwhelmingly likely that ETIs exist, but that it's also overwhelmingly *un*likely that we will ever be able to make contact with them. If this were to be the case, it would make no theoretical or practical difference whether some of these ETIs happen to exist contempor-

[4]For a history of the Drake equation, see Dick (1996).

aneously with our own civilization. The bare fact of their existence, anywhere and at any time, would make no less of an impact on our lives and thoughts than the fact of their existence now. Thus it's more perspicuous to equate the ETI hypothesis with the general proposition that there are ETIs somewhere and at some time. This means that the Drake equation relevant to the ETI hypothesis is at least two terms shorter than the equation relevant to the SETI hypothesis. Let N' be the number of planets in the galaxy harboring intelligent life at any time. Then

$$N' = N_* f_p n_e f_l f_i.$$

When it comes to SETI, however, even the full equation for N given above is two terms short. It's not enough to know that there exist temporally overlapping extraterrestrial civilizations engaged in interstellar communication. Conceivably, we might come to be persuaded that this is so on the basis of a theoretical argument. But this would not yet be SETI success. To succeed at SETI, we have to *detect their signals*. Moreover, having detected their signals, we need to ascertain that these signals are *expressions in a language* rather than a peculiar inanimate phenomenon, or the radio equivalent of a bird song. The most straightforward way to identify a series of signals as linguistic is to *decode* it—to figure out a scheme for translating the signals into our language and to test our scheme on a new series of signals. Now Donald Davidson (1973–4) has famously expressed the view that decoding is the *only* way to ascertain that a series of signals belongs to a language. I won't assess the merits and demerits of this view in this paper. But it's worth noting that if Davidson is wrong, then *CETI* (not SETI) success requires a third extra-Drakean step beyond detecting extraterrestrial signals and ascertaining that they're linguistic: it requires also that these linguistic signals be decoded. If, on the other hand, Davidson is right, then the second step is identical to the third, and SETI and CETI succeed or fail together.

As would be expected, the Drake equation incorporates factors from all the major branches of science—the physical, the biological and the psychosocial. Most of the discussion among SETI researchers has centered around the physical and biological factors, represented at the beginning and middle of the equation. The psychosocial factors at the end have not usually been the loci of sustained intellectual dispute. What accounts for this disparity? I attribute it to the fact that SETI funding has gone almost exclusively to physical and biological scientists.[5] An alternative explanation is that there isn't as much to be said about the psychosocial issues. There are no canonical psychosocial theories that can be brought to bear on the issues. Physics and biology, on the other hand, possess the heavy intellectual artillery that must surely make a difference in the field. This sounds very plausible. However,

[5]Dick's historical summary of SETI project funding makes no mention of any grants for the investigation of psychosocial factors. See Dick (1996), pp. 494–501.

more than one critic has noted that our physical and biological theories, excellent as they are in their own right, fail to make contact with the task of estimating the Drake factors (McMullin, 1980; Baird, 1987; Mash, 1993). The paths from current physical and biological theory to estimates of the Drake factors are all paved with adhockery. In the end, the physical and biological parameters of the Drake equation are in the same epistemic boat as the psychosocial: our estimates of them are not dictated by theoretical principles.

McMullin draws a harsh lesson from this state of affairs:

> Until we have a theory of some sort, we cannot attach a theoretical probability of any kind to an outcome . . . there is no responsible way of [estimating the Drake factors] except in the light of a detailed and properly warranted theory. (McMullin, 1980, p. 84)

McMullin's judgment may be too severe, for the likelihood of a hypothesis can also be estimated on the basis of brute enumerative induction, even in the absence of a detailed relevant theory. If numerous randomly chosen samples of the raven population turn out to be composed of about 60% black ravens and 40% non-black ravens, it would be rational for us to ascribe a probability of 0.6 to the hypothesis that the next raven we encounter will be black—even if we lack a theoretical account of why the observed frequencies are what they are.

SETI skeptics point out, however, that this possibility will not help us to develop rationally grounded opinions about the Drake factors. The reason is simple: you can't base an inductive inference on a single case. If we had observed only a single raven and found it to be black, we would have no more inductive warrant for expecting future ravens to be black than for expecting that other ravens will have precisely the same scratch on their left foot as one that was found on the original, observed raven. With just a single case before us, there's no basis for making an educated guess as to which of the observed properties are connected in some law-like way to ravenhood, and which are merely accidents. The fact that we've seen one inhabitable planet spawn life doesn't give us inductive grounds for believing that any of the other inhabitable planets in the galaxy will also spawn life. Nor is a single case sufficient for making a probabilistic estimate. In the absence of relevant theory, the only way to make such an estimate is to equate the probability of the target event to the frequency of target events in the total sample of observed events. If we applied this procedure to the Drake factors, we would conclude that the probability that a randomly chosen star has a planet which harbors an advanced civilization is *one*, since the observed frequency of such stars in the set of adequately examined stars is one over one. The catch, of course, is that the reliability of the probabilistic estimate increases with the size of the observed sample. When the sample contains a single element, the estimate is maximally unreliable.

So the single case of intelligent life that we're familiar with doesn't underwrite any estimate for the probability that there are other cases. However, it's widely

believed in SETI circles that (1) the single case does allow us to conclude that the probability of there being intelligent life on another, randomly chosen world cannot be exactly zero, and that (2) this is ultimately enough of a concession to warrant acceptance of the hypothesis that there *is* intelligent life on *some* other worlds. Let's look at both parts of this package. The rationale for the first part runs as follows: the single case of intelligent life on earth establishes that the evolution of intelligent beings is not forbidden by the laws of nature, and what is not forbidden has a non-zero probability of recurring. There are hidden complexities in this inferential step that have never been adequately analyzed. The main problem with it is that the occurrence of an event doesn't logically entail that the laws that produced it must allow for a second production. It's easy to construct logically possible universes whose laws entail that certain events are unique. For example, consider a universe that begins to exist at time t_0 containing only a single particle of size S. Suppose also that the only law of nature in this universe is that after a given time interval, every particle in the universe emits an 'offspring' particle that's half the size of the original. Thus after a single time interval, there are two particles: the original parent particle of size S, and an offspring of size S/2. After the second time interval, there are four particles: one of size S, two of size S/2, and one of size S/4; and so on. This is a well-behaved logically possible universe. But in this universe, it would be a mistake to suppose that just because there exists one instance of a particular type of item, there must be a non-zero probability that there are other instances of the same item. In this universe, there exists a particle of size S—namely, the original parent particle. Nevertheless, the probability of encountering *another* particle of size S is exactly zero. Additional items of size S have the same status in this universe as perpetual motion machines have in ours: they're ruled out by natural law. The lesson is that the assumption that the universe is a lawful place doesn't by itself underwrite the conclusion that what happens once must be allowed to happen again. It must be regarded as an open question whether the laws of *our* universe entail this conclusion.

Moreover, the inference from our single case to the non-zero probability of ETI wouldn't do ETI advocates any good, even if they were permitted to make it. There are numerous arguments in the ETI literature that attempt to parlay the non-zeroness of the probability of ETI into a reason for accepting ETI as a plausible hypothesis. All of these arguments, however, commit one or both of two errors: (1) the error of one-sided estimation, and (2) the error of over-precision. The main example of the first error is to be found in the argument that the universe is so vast that even extremely unlikely events are likely to occur at some place or other:

> [T]he idea that *we* are the only intelligent creatures in a cosmos of a hundred million galaxies is so preposterous that there are very few astronomers today who would take it seriously. (Clarke, 1973, p. 90)

> With the Universe constructed on so vast a scale, it would seem inherently improbable that our small Earth can be the only home of life. (Jones, 1940, p. 7)

> In fact, there are billions of galaxies, each containing billions of stars. This is the single most important reason for optimism in the search for life. In a universe so vast, with stars as numerous as grains of sand, it is hard to imagine that the conditions for life have not arisen elsewhere. (McDonough, 1987, p. 59)

To call this an argument is to insult arguments. In fact, it's nothing more than a profession of one's faith that the number of stars is large enough to overwhelm the smallness of the probability that there are ETIs associated with a randomly selected star. Certainly this conclusion doesn't follow from the admission that the probability of ETI is non-zero. For however many stars there may be, so long as the number is finite, there's going to be a non-zero probability for intelligent life that's so ultramicroscopically small that the net probability of ETI anywhere in the universe will still be as close to zero as makes no difference. For example, the estimated number of stars in the observable universe is 10^{22}. Let's suppose that this is a drastic underestimation. Let's suppose that there are actually 10^{100} stars. Let's also suppose that all the Drake factors except one are extremely favorable for the SETI hypothesis: that virtually all stars have planets, that the average number of planets per solar system whose environment is suitable for life is 100 (we live in an extremely deprived system), that all planets suitable for life actually develop life, that advanced civilizations arise every place where intelligence develops, and that once an advanced civilization arises it lasts forever. The only fly in the ointment is that the probability that intelligence develops where life exists is only 10^{-150}. Then, even though there are 10^{102} planets harboring life, the probability that any of them contains *intelligent* life is only 10^{-48}. If that isn't negligible enough to eliminate ETI from serious consideration, we can run through the same example with a probability of 10^{-1000} for the development of intelligence, or $10^{-1,000,000}$. The point is that the thesis that the probability of ETI around a given star is nonzero is compatible with any of these estimates.

That's all that needs to be said about the raw appeal to large numbers. But Mash (1993) has an exercise that helps to dissolve the unwarranted intuition that the immensity of the universe must overwhelm the minuteness of the probability of ETI in any given locale. The probability of passing the Drake gateways gets smaller and smaller as more of them are taken into account: it's likelier that there are inhabitable planets than that there are inhabitable planets that actually possess life; it's likelier that there's extraterrestrial life than that there's extraterrestrial *intelligent* life; and so on. The Drake equation stops at intelligent beings who are sufficiently like us that they engage in their own SETI project. The pro-SETI intuition is that, though the probability gets smaller as we traverse the various Drake factors, we're left at the end with a probability that's large enough to be swamped by the vastness of the universe. Now what happens if we add an additional requirement to the end of the Drake equation—say that the ETIs not only engage in interstellar communication, but that they also speak English? To be sure, the resulting scenario would be orders of magnitude less likely than the scenario that merely specifies

that the ETIs engage in interstellar communication. But why not say, once again, that the vastness of the universe ensures that even these minute probabilities are actualized? After all, it would be utterly arbitrary and unreasonable to suppose that net probabilities cease to be overwhelmed by the vastness of the universe at the precise point where the conventional Drake equation leaves off—that it's plausible that there are ETIs who build radio telescopes somewhere in the universe, but that it isn't plausible that there are English-speaking ETIs who build radio telescopes.

And why stop at speaking English? Why shouldn't we suppose that the vastness of the universe overwhelms the extreme unlikelihood of there being ETIs who have come upon Velcro fasteners, and whose teenagers amuse themselves with skateboards, hula-hoops, and yo-yos? In fact, why doesn't the vastness of the universe make it plausible that there's a twin earth where someone just like you reads these very words? The argument is the same: it's scientifically possible for there to be a twin earth—or so we have assumed; therefore, given the billions upon billions of stars that each provide an occasion for it, it's likely to have happened somewhere. But surely there's some point along the chain of diminishing probabilities where even the most ardent SETI enthusiast will balk—if not at the existence of one twin earth, then at the existence of ten, or a million of them. However large the universe may be (so long as it's finite), it isn't going to be vast enough to accommodate any and all scientifically possible scenarios. There has to be a cut-off point along the series of diminishing probabilities beyond which the vastness argument no longer applies. But then, what reason do we have for supposing that intelligent life lies on one side of this cut-off point, while the employment of Velcro fasteners and skateboards lies on the other? Rescher has put the matter very succinctly:

> Admittedly, cosmic locales are very numerous. But probabilities can get to be very small: no matter how massive N, there is that diminutive 1/N that can countervail against it. (Rescher, 1985, p. 109)

It's important to realize that this argument cuts both ways. Mash reminds us that

> the corollary to [Rescher's assertion] is that no matter how small a fraction 1/N faces us, there is always that big number N that can squash it. (Mash, 1993, p. 220)

SETI skeptics have tended to make the symmetric error of arguing one-sidedly for the rejection of ETI on the grounds that 1/N is very, very small. An example is Mayr's 'evolutionary' argument:

> The point I am making is the incredible improbability of genuine intelligence emerging [from an evolutionary process]. There were probably more than a billion species of animals on earth, belonging to many millions of separate phyletic lines, all living on this planet earth which is hospitable to intelligence, and yet only a single one of them succeeded in producing intelligence. (Mayr, 1985, p. 28)

At first glance, it looks as though this anti-SETI evolutionary argument is superior to the pro-SETI vastness-of-the-universe argument we've just considered: the latter

employs an inductive inference based on a single case, whereas the former's inductive inference is based on billions of cases. Of course, the billions of cases aren't randomly selected. They're all species that evolved on the same planet. Mayr characterizes the earth as a planet 'which is hospitable to intelligence'. But this is nothing more than a guess. It could turn out that there's something about the terrestrial environment that *inhibits* the evolution of intelligence, as compared to other planetary environments which are suitable for life. Alternatively, there might be something about evolutionary processes generally that militates against the appearance of more than one intelligent species per planet. (Perhaps intelligence is inevitably accompanied by a xenophobia so intense that the first intelligent species to appear exterminates all the near-intelligent competitors.)

An even more telling critique of the evolutionary argument is that it's an induction based on a single case after all. The one-in-billions figure alluded to by Mayr is the observed frequency of intelligent species among total species. But this frequency does not provide an estimate of any of the factors in the Drake equation. The factor that comes closest to it is f_i, the fraction of life-bearing planets on which intelligence evolves. But the appropriate ratio for estimating this factor is the observed incidences of intelligent species among total *evolutionary histories*. It makes no difference to f_i if evolutionary processes profligately produce myriads of unintelligent species. All that matters is how frequently evolution produces intelligence. But of course, we have only one evolutionary history upon which to base our opinion. And *that* observed frequency is one over one.

But let's suppose that the contemplation of the twists and turns of evolutionary history on earth persuades us all that the probability that intelligence evolves on a life-bearing planet is extraordinarily small. The question once again arises: is it small enough? The anti-SETI biologists seem to be guilty of relying on an unreasoned faith which is complementary to the unreasoned faith of the pro-SETI astronomers. In the conflict between enormous Ns and minute 1/Ns, the astronomers arbitrarily take the side of the Ns and the biologists arbitrarily align themselves with the 1/Ns. One-sided appeals to enormity or minuteness are not enough. To ground an opinion about SETI, we need to make absolute estimates of *all* the Drake factors, multiply them together, and see what we get. Such estimates are available in the SETI literature (Dole, 1964; Sagan, 1980). But in the absence of both a theoretical and an inductive basis for these estimates, how can the estimators defend themselves against the charge that they're pulling their numbers out of a hat?

There's a line of defense for Drake estimators that isn't discussed in Mash's otherwise comprehensive analysis. Why not perform an induction over the educated guesses of successful scientists? The idea is that if the past guesses of a population of persons or devices have proven to be right more often than can be accounted for by chance, we have inductive grounds for accepting subsequent guesses from members of the same population. It's arguable that success in the empirical sciences

requires an intuitive ability to make good guesses on the basis of unsystematic, holistic considerations—or else the scientist would not be guided to undertake what turns out to be a fruitful research program. On this account, successful scientists are conceived to be oracular devices for whose accuracy we have good inductive evidence, even though we lack an account of how they arrive at their estimates. If an eminent astrophysicist guesses that the fraction of stars possessing planets is about 1/10, that's at least *some* evidence that the fraction is about 1/10.

One problem with this induction over successful scientists is that the successful scientists are not of one mind about ETI. There is a pattern, though: the physical scientists (such as Sagan) tend to be pro-ETI, while the biologists (such as Mayr) tend to be anti-ETI. The indicated procedure is obvious. Successful physical scientists have proven themselves to be good estimators specifically of physical factors, and successful biologists have proven themselves to be good estimators of biological factors. The logic of induction dictates that we should accept the estimates made by each group of the factors relevant to its own area of expertise—the physical scientists' estimates of the physical factors and the biologists' estimates of the biological factors.

Here's why this mode of resolution won't work. Suppose I'm asked to estimate the probability that it will rain on a randomly selected day in Kathmandu. I have no theoretical knowledge of meteorology that bears on this question, nor do I have any detailed frequency information that would enable me to tackle the question inductively. Still, I can come up with a ball-park figure. I remember that Kathmandu is not a desert, so the probability isn't anything as low as 0.01. In fact, I think it's reputed to have a rainy season of several months during which it rains virtually every day. So the probability of rain on a random day is going to be at least 0.2. By relying on vague memories and anecdotal reports of this type, I can come up with an estimate that has some rational foundation. At least it would be more rational to accept the estimate than to accept a figure for the probability that's literally drawn out of a hat. It would be exceedingly bizarre, however, if my ruminations based on vague memories and unsystematic anecdotes led me to the estimate that the probability of rain in Kathmandu is 0.3987. It would be even more bizarre if another person, to whom I had conveyed this estimate, tried to get me to change it to 0.3986. It's utterly implausible to suppose that the sets of memories and anecdotes that warrant an estimate of 0.3987 are discriminably different from the sets of memories and anecdotes that warrant an estimate of 0.3986. To make either estimate is to commit the *error of over-precise estimation*—i.e., to specify a quantity with a degree of precision which is unwarranted by the means available for its measurement or estimation.

Now we've seen that the resolution of the ETI question depends on the outcome of a struggle between very large numbers and very small numbers. Suppose that it's been estimated that there are 10^{22} inhabitable planets in the universe. (This is more or less in line with the estimates of pro-SETI physical scientists.) Then if

the frequency of life on an inhabitable planet and of intelligence on a planet containing life are both one in a million, the estimated number of planets with intelligent life would be an enormous 10^{10}, and the rational acceptability of the ETI hypothesis would be assured. But if both those frequencies were taken to be one in a trillion, the estimated number of planets would be zero (to the nearest integer). So the disposition of the ETI hypothesis may very well turn on whether some frequency estimates are one in a million or one in a trillion.

How reasonable is it to suppose that the educated guesses of successful scientists are sufficiently refined to discriminate between a frequency of one in a million and a frequency of one in a trillion? From one perspective, the difference between the two frequencies is very large. The fair value of a lottery ticket which gives us one chance in a million to win is one million times the fair value of a ticket where the chance of winning is one in a trillion. If the prize is a million dollars, the value of the former ticket is one dollar, while the value of the latter ticket is 1/10,000 of one cent. From another perspective, however, the difference between these two frequencies is very small. In fact, the amount by which one frequency exceeds the other is substantially less than the difference between the two estimates of 0.3986 and 0.3987 for the probability of rain in Kathmandu. So is the difference between one in a million and one in a trillion very large or very small? The answer depends on one's purpose. If the purpose is to calculate the fair value of lottery tickets, the difference is very large. The purpose here, however, is to determine the ease with which the two frequencies can be *discriminated* on the basis of holistic and unsystematic informational input. For *that* purpose, it seems compelling that the difference between one in a million and one in a trillion is fantastically small. It may be reasonable for us to accept successful scientists' guesses that some particular Drake frequencies are either 'large' or 'moderate' or 'small'. Perhaps it's even reasonable to suppose that they can tell us whether a frequency is merely 'small' or 'very small' or 'very, very small'. But it's inconceivable that there could be two sets of discriminable holistic tableaus, one of which would warrant the one-in-a-million estimate and the other the one-in-a-trillion estimate. Such a feat of discrimination would be akin to discriminating visually between a million-sided polygon and a trillion-sided polygon. To insist on either value is to commit the error of over-precise estimation. Yet the two values quite likely give different answers to the ETI question. Therefore we can't expect to obtain a rationally compelling answer to the ETI question based on the educated guesses of successful scientists.

3. The Lucretian Argument

What about the Lucretian move from mere enormity to literal infinity? While it's true that no finite N is big enough to do the job that's required of it in the one-sided vastness-of-the-universe argument, the situation is changed qualitatively if it's supposed that there are infinitely many occasions for an event. The argument, in a nutshell, is that if the number of occasions for an event is infinite, then no

matter how close to zero the probability of its occurrence may be, it's bound to happen. A monkey striking typewriter keys at random forever will eventually type the verbatim text of *Hamlet*. By the same token, if the world is infinite in space or time, isn't it equally inevitable that the random collision of atoms will produce another intelligent species? The modern version of the Lucretian argument comes to an even more radical conclusion than Lucretius. The latter merely claimed that the falsehood of the ETI hypothesis was 'in the highest degree unlikely'. The contemporary claim is that if the universe is infinite, then there are *bound to be infinitely many worlds with ETIs*:

> Analysis shows that . . . an open universe must be infinite in extent, with an infinite number of galaxies, an infinite number of stars, and an infinite number of planets. In an infinite universe, any event which has a finite probability—no matter how small—of occurring on a single given planet must inevitably occur on some planet. In fact, such an event must occur on an infinite number of planets. (Hart, 1982, p. 163)

> If cosmological initial conditions are exhaustively random and infinite then anything that can occur with non-vanishing probability will occur somewhere; in fact, it will occur infinitely often. (Barrow and Tipler, 1988, p. 7)

These opinions are based on the firm ground of modern probability theory—specifically, the Borel-Cantelli theorem, which specifies that in an infinite number of independent trials, any event with non-zero probability will occur infinitely often (Chow and Teicher, 1978). It's important to note, however, that the Borel-Cantelli theorem doesn't quite underwrite the existence of ETIs in an infinite universe. The firmly grounded thesis is that there are surely ETIs in an infinite universe *if* the probability of ETI is non-zero. But how do we know that the laws of the universe allow for ETIs? The pro-ETI reply is that this antecedent condition is assured by our own existence: the fact that there's intelligent life on earth establishes that the probability of intelligent life developing on a randomly selected planet is greater than zero. This argument is too quick, however. It overlooks the possibility that the laws of nature may specify that certain events occur exactly once—i.e., that the probability of some *recurrences* may be exactly zero. This issue came up in the previous section, where a well-behaved universe was described whose laws did in fact entail that certain events were unique. Though it wasn't relevant to mention it at the time, this universe also happened to be temporally infinite. Thus it isn't necessarily true that every event that occurs once must inevitably repeat itself in an infinite universe.

Of course, the universe described in the previous section was an exceedingly implausible place. Perhaps if we restrict ourselves to universes that satisfy certain plausible conditions, we will find that in all *these* places, what occurs once must have a non-zero probability of recurring. Mash (1993) lists the following six conditions that must be satisfied if the Lucretian argument is to go through.

1. The first condition is, of course, that the universe must be infinite. It makes

no difference whether it's infinite in space or in time. One monkey randomly typing forever will eventually produce *Hamlet*, and infinitely many monkeys randomly typing for a single day will produce *Hamlet* on that day.

2. The number of different kinds of materials in the universe must be finite. 'If the keyboard of the Shakespearean typewriter contains an infinite . . . variety of types of keys, recurring sequences need not occur' (Mash, 1993, p. 204).
3. The state of affairs to be repeated must itself be finite. If *Hamlet* were infinitely long, then even a monkey typing forever might never produce it.
4. The laws of nature must be the same everywhere in the universe. Otherwise, the universe may be partitionable into infinitely many finite-sized sub-universes, each of which is governed by different laws. By virtue of its finite size, one of these sub-universes may contain an event which is unique to that sub-universe. By virtue of their operating under different laws of nature, that unique event may be forbidden to occur in the other sub-universes. The result would be an event which occurs only once in the whole universe.
5. All objects must be aggregates of specifiable, discrete elements. It seems to me that this condition of Mash's is tautologously satisfied in every logically possible world. For suppose that some type of object X is *not* describable as an aggregate of other elements. Then X is itself an element. At the very least, every object is an aggregate composed entirely of itself. It's true that the Lucretian doesn't want to add infinitely many types of indecomposable objects to the class of elements, for the existence of infinite varieties means that recurrence isn't assured after all. But this is ruled out by condition 2.
6. Determinism is false. If some processes in the world weren't inherently stochastic, then ETI might be ruled out by deterministic laws even in an infinite universe. 'If determinism is thoroughgoing, the typewriter might as easily produce an infinite sequence of 'a's as all possible sequences' (p. 210).

All of Mash's conditions correspond at least to scientifically respectable hypotheses. The most dubious one is the hypothesis that the world is in fact infinite. But even this is by no means definitively ruled out by modern cosmology. So can ETI advocates claim that, given our current scientific understanding, ETI is a sure thing *if* the universe is infinite? The answer is: not yet. Mash shows only that his conditions are *necessary* for the Lucretian argument to go through. He doesn't attempt to show that they're *sufficient*. In fact, they surely aren't sufficient. It's trivially easy to construct universes that satisfy all six Mash conditions, but which contain events that occur only once. The hypothetical universe described in Section 2 won't do the job because it's a deterministic world, in violation of condition 6. But consider a world which is infinite in space and time (condition 1), and contains nothing other than a finite number of different types of discrete elements (conditions 2 and 5) that move about in accordance with inherently probabilistic laws (condition 6) which are the same everywhere (condition 4). Suppose also that the elements are indestructible and uncreatable: their number is, always has been, and always will

be the same. Now there's nothing in Mash's conditions that requires that there be infinitely many particles of each type. Suppose then that there are exactly n particles of type T, where n is a finite number. Define a T_n-object as any aggregation of n T-type particles, irrespective of the distances between the individual particles. This object is finite, which satisfies the last remaining Mash condition (condition 3). Yet it seems that there exists, always has existed, and always will exist exactly one T_n-object in the universe. The fact that one of them exists is evidently not a reason for supposing that there's a non-zero probability that another one exists somewhere else. If the reader has qualms about accepting T_n-objects as bona fide objects, then consider a universe just like the one described above, except that it contains only a single T-type particle. Then that particle is eternally unique.

I don't, of course, claim any plausibility for this artificial world, nor do I claim that ETI would be any more problematic in it than in some other imaginable world. The point of the construction is only to show that Mash hasn't yet formulated a set of premises which is sufficient to bear the Lucretian argument to its conclusion. At present, nobody knows exactly what a sufficient set of premises is going to look like. For all we know, the list will include conditions that our universe clearly fails to satisfy. Until we know otherwise, Lucretian considerations don't provide us with rational grounds for believing in ETI, nor even for believing in the conditional proposition that there are ETIs if the universe is infinite.

Moreover, when it comes to the *SETI* hypothesis that we might be able to acquire direct empirical evidence of the existence of ETIs, the infinity of the universe becomes irrelevant. According to relativity theory, we can receive signals from only a circumscribed region of space-time—the region from which signals travelling at the speed of light are close enough to have arrived. If the universe has an infinite past, this 'light cone' could theoretically encompass an infinitely large portion of space-time. But here we must reckon with another canonical theory of modern cosmology—the Big Bang theory. The Big Bang theory doesn't logically entail that the universe has a finite past. It's compatible with the theory, for instance, that the Big Bang was preceded by an era in which the universe contracted to a point of maximal density—a Big Crunch. But if this is so, then the Big Crunch poses an insurmountable informational barrier. It would be impossible for any empirical information about pre-Big-Crunch conditions to survive the crunch. So one way or another, our contemporary scientific views entail that the *observable* universe is a finite chunk of space-time. Therefore Lucretian arguments, even if they turn out to be sound, provide no support for SETI.

4. The Argument from Mediocrity and Indifference

There's one more pro-ETI argument to evaluate—the argument based on the *principle of mediocrity*. Despite its frequent employment, no precise statement or adequate analysis of this principle exists. The principle seems to be that we should take it as a default assumption that the object or event under consideration is more

likely to be average or typical of its type than unusual or unique. In the absence of compelling reasons to suppose otherwise, our best guess is supposed to be that 'our surroundings are more or less typical of any other region in the universe' (Shklovskii and Sagan, 1966, p. 357), that 'there is nothing special about the solar system and the planet earth' (Rood and Trefil, 1981, p. 4), that 'our own evolution is typical' (Tipler, 1985, p. 134), and so on. The persistent application of the principle magically transforms our ignorance into a case for the existence of ETIs:

> In estimating n_e, for example, we start by allowing (on the basis of the assumption of mediocrity) that the environmental conditions on many other planets are roughly like those of the Earth four billion years ago. Then we take the next term, f_l, and argue (on the basis of something like Miller and Urey's 1959 'primordial soup' theory) that in such conditions life will likely arise. But now if our theory on this point is incomplete, imprecise, or found wanting in other respects, we may leap the hurdle by simply granting ourselves (again, on the basis of the assumption of mediocrity) that life is not uncommon on planets with Earthlike conditions, and then proceed to f_i, arguing from the basis of natural selection that intelligent life will arise—and so on through the equation. The assumption of mediocrity thus acts as a kind of fall-back, a temporary (we hope) stopgap that lets us get on with the theory at hand. (Mash, 1993, p. 214)

One has to admire the sheer audacity of this grab for credibility. Having failed to arrive at a substantial probability for ETI by appeals to the vastness of the universe, ETI advocates now claim to have achieved their goal *without* relying on the vastness of the universe.

What, exactly, does the principle of mediocrity tell us about probabilities? Here's a first pass at an interpretation: if an item is drawn at random from one of several sets or categories, it's likelier to come from the most numerous category than from any one of the less numerous categories. In the case of ETI, the single case of planet earth is drawn from the category of 'inhabitable planets containing life', as opposed to the category 'inhabitable planet not containing life'. Therefore, since the principle of mediocrity tells us that it's likelier to have been drawn from the more numerous category, we conclude that the set of inhabitable planets containing life—the set from which the earth was in fact drawn—is probably more numerous than the set of inhabitable planets not containing life. In point of fact, SETI advocates don't actually draw the inference that the set of life-bearing planets is *more* numerous than the set of non-life-bearing planets. The conclusion is the rather vague claim that the proportion of life-bearing planets is not inconsiderable. Presumably, the principle of mediocrity ascribes progressively smaller probabilities to ever-smaller estimates of the proportion of life-bearing planets, the least likely scenario being that our single case is unique. The criticisms to be wielded below against mediocrity reasoning are applicable, *mutatis mutandis*, to the vaguer and more convoluted principle. But the critique itself would become convoluted, and would make for an unpleasant stretch of reading. So I'll run my discussion on the original, simplified formulation—the one that says that the set from which the single sample is selected is more likely than not to be more numerous than its

complement. Everything that's wrong with this claim is also wrong with the vaguer claim.

There are two problems with the ETI argument from mediocrity, both of them utterly debilitating. The first is that whatever prima facie plausibility the principle of mediocrity may have is entirely dependent on the single case having been drawn at random. But the earth is not a randomly selected planet:

> . . . we did not 'choose' the Earth randomly. It is trivially true that the only planets on which intelligent conscious life like ours can evolve are those on which the conditions for that evolution are met. It is not as though we were like some extraplanetary being walking in on the galaxy from the outside . . . and choosing a planet at random. We have no choice but to start with a planet on which intelligent life has evolved, but that in no way insures that the case is typical. (Apparently ETI proponents would like to argue that if, e.g., habitable planets were very, very rare, we would have found ourselves inhabiting a 'typical' uninhabitable planet rather than the Earth.) So even if the assumption of mediocrity is a valid means for licensing induction from a single case, it may not be applicable for certain crucial inferences along the Drake equation. (Mash, 1993, p. 218)

The problem of randomness aside, the principle of mediocrity is demonstrably defective. Recall its formulation: if an item is drawn at random from one of several sets or categories, it's likelier to come from the most numerous category than from any one of the less numerous categories. This statement is amenable to two drastically different readings, one of which is a probabilistic truism, the other a fallacy. The principle that's needed to underwrite ETI is the fallacious version. But the fallacy is obscured by virtue of its being confused with the truism. On one reading, the principle states that the single randomly drawn object is more likely to have come from the category that we know to be more numerous. This is the truism. If category A contains 3 elements and category B contains 1 element, then a random draw from the total population of 4 elements has a 3/4 probability of having come from A, and only a 1/4 probability of having come from B. This inference presupposes that we have antecedent knowledge of the relative numerosities of the classes A and B. In its ETI application, however, our antecedent knowledge and the inference we draw from it are reversed. We know that the random choice has come from A, and we infer from this that A is probably more numerous than B. For example, the classes A and B are 'inhabitable planets that contain life' and 'inhabitable planets that do not contain life', respectively, and the fact that our single examined case belongs to A is alleged to license the inference that A is probably more numerous than B (more vaguely, that the proportion of As is not inconsiderable). This is an altogether more speculative inference than the first.

Mash provides an indirect refutation of the mediocrity assumption—that is, he shows that it entails consequences which even the most sanguine SETI advocate would be loath to accept. Suppose that our single case is evidence for the greater numerosity of the class to which it belongs than of its complement. Then why should we only apply it to the categories which are represented in the Drake equ-

ation? Once again, our planet contains Velcro-clad, skateboarding teenagers. By parity of reasoning, our planet must be regarded as probably average or typical with regard to the incidence of Velcro-clad, skateboarding teenagers. Therefore, it's likely that Velcro-clad, skateboarding teenagers reside on most (more vaguely: many) of the inhabitable planets in the universe. The argument is the same. Indeed, the argument is the same all the way to twin Earths. Since its unconstrained use leads to absurdity, there must be some additional condition for the applicability of the mediocrity principle. But there's no hint of what this additional condition might be in the ETI literature.

Mash's argument is an entirely adequate refutation of the mediocrity principle. But it doesn't provide an analysis of where the reasoning goes wrong. The reasoning goes wrong, I think, by virtue of a tacit appeal to the notorious probabilistic *principle of indifference*, which is known to lead to contradictory conclusions. The conclusion of the mediocrity principle is that A, the category from which our single case is known to come, is *probably* more numerous than its complement B. This probabilistic claim means that, in light of the knowledge that the single examined case comes from A, there are more ways for A to be more numerous than B than there are ways for B to be more numerous than A, relative to some equiprobable set of 'ways'. The source of the temptation to suppose that this might be so is highlighted in the following minimal scenario. Suppose we know that there are three inhabitable planets in our stellar system, including our own. We know that our own planet contains life, but we have no evidence relating to the existence or non-existence of life on the other two. Given only this information, is it any likelier that there are more life-bearing planets in our system than non-life-bearing planets? If the principle of mediocrity is valid on a galactic scale, it should also be applicable to this local planetary scenario. Our own planet would be at least marginally more typical if at least one of the other planets contained life than if neither of them did. Therefore, if the mediocrity principle is valid, it should be at least marginally more probable that one or both of the other planets contains life than that neither of them does. Is this a valid probabilistic inference?

The argument that we're tempted to make is the following. The two unexamined planets—call them P and Q—may each either harbor life or not harbor life. The four possibilities are therefore: (1) neither P nor Q harbors life, (2) both P and Q harbor life, (3) P harbors life but Q doesn't, and (4) Q harbors life but P doesn't. If (1) obtains, there is only one life-bearing planet (our own), as compared to two non-life-bearing planets. If (2) obtains, there are three life-bearing planets (our own included) and no non-life-bearing planets. If either (3) or (4) obtains, there are two life-bearing planets (ours and one other) and one non-life-bearing planet. In all of these scenarios except the first, the number of life-bearing planets exceeds the number of non-life-bearing planets. Therefore, the probability that the set of life-bearing planets is greater than the set of non-life-bearing planets is 3/4.

This conclusion obviously depends on the assumption that the four scenarios are

equiprobable. Why should we make that assumption? Well, we've assumed that we have no knowledge relevant to the question whether the unexamined planets contain life. So the judgment that the four scenarios are equiprobable can't be based on any relevant evidence. It must, in fact, be based on our very lack of knowledge. There is no other candidate. We must be assuming that in the absence of any reason for preferring one hypothesis over another, we should regard them as equiprobable. This is the principle of indifference.

Unfortunately for ETI and SETI, the principle of indifference is invalid. The problems of indifference reasoning are discussed extensively by Salmon (1966). The root problem is that being in a state of epistemic indifference between complementary hypotheses A and B is entirely compatible with also being indifferent with regard to hypotheses A, B_1 and B_2, where B_1 and B_2 are an exhaustive decomposition of the state of affairs indicated by B. The indifference principle thus stipulates that the probability of A is 1/2 by virtue of the toss-up between A and B, and that it's 1/3 by virtue of the toss-up between A, B_1, and B_2. More generally, the existence of a state of indifference between A and B is compatible with the existence of states of indifference between A, B_1, B_2, . . ., and B_n, for any n. For example, suppose that I think I have no basis whatever for deciding whether any gods exist, or if any do exist, what their number might be. To avoid having to deal with inessential complications, suppose that this uncertainty over the number of gods extends from zero to 9: I'm sure that the number of gods is in the single digits, but I have no idea what that digit might be. This is surely a logically possible epistemic state. What does the indifference principle tell us about this state? I'm indifferent between the hypothesis that there are no gods and the hypothesis that there are some gods. Therefore the probability that there are no gods is 1/2. But I'm also indifferent between the ten hypotheses of the form 'There exist(s) exactly n god(s)', where n runs from 0 to 9. Therefore the probability that there are no gods is 1/10. Contradiction.

It's worth noting that the conjunction of the two states of indifference that lead to conflicting probabilities is sometimes more than merely possible: *being in one of the states sometimes logically entails being in the other*. This is the case in the foregoing example. For suppose I encounter a bit of evidence which asymmetrically supports the hypothesis that there are no gods over the hypothesis that there are some gods. Then it follows that I'm also no longer indifferent between the ten hypotheses of the form 'There exist(s) exactly n god(s)'—for my single bit of evidence favors the hypothesis with $n = 0$ over each of the other 9 hypotheses. Thus indifference over the ten numbered hypotheses *entails* indifference between existence and non-existence. The contradictory conclusions to be drawn by the indifference principle are not only available—they're required. The fact that some states of indifference entail other indifferences that lead to contradictory probabilities is well known in the literature of probability theory, where it goes by the name of the Bertrand paradox (von Mises, 1957, p. 77). In retrospect, it's not

surprising that indifference reasoning would turn out to be untenable—it's just too good to be true to suppose that complete ignorance can be transformed into precise probabilistic information:

> Knowledge of probabilities is concrete knowledge about occurrences; otherwise, it is useless for prediction and action. According to the principle of indifference, this kind of knowledge can result immediately from our ignorance of reasons to regard one occurrence as more probable than another. This is epistemological magic. (Salmon, 1966, p. 66)

Let's return to the scenario of the three planets. It's assumed that we have complete uncertainty regarding whether planets P and Q contain life. It follows that we're indifferent as to the four indicated states of affairs: (1) P and Q contain life, (2) neither P nor Q contains life, (3) P contains life but Q doesn't, and (4) Q contains life but P doesn't. Applying the principle of indifference, we obtain a probability of 3/4 for the hypothesis that there are more life-bearing than non-life-bearing planets. But our complete uncertainty also entails that we're indifferent as to the hypotheses (1) that there is life on other planets and (2) that there is no life on other planets. Applying the principle of indifference to *this* decomposition of the possibilities, we obtain a probability of 1/2 for the hypothesis that there is life on other planets. But the 1/2-probability hypothesis that there is life on other planets is identical to the previous 3/4-probability hypothesis that there are more life-bearing than non-life-bearing planets, which is impossible. The problem with the principle of indifference isn't merely that it gives an *erroneous* probability estimate—it's that it doesn't give any determinate estimate at all.

Moreover, if we're going to permit ourselves to use the principle of indifference, then mediocrity arguments—considerations of our 'typicalness'—are supernumerary. We could simply plead complete ignorance of any factor that might bear on the truth or falsehood of the ETI hypothesis, including whether or not our planet, our evolution, etc., are typical or unique. Since there's no reason to favor either the ETI hypothesis or its negation, we could claim that the principle of indifference gives us a probability of 1/2 for each one—certainly high enough to warrant a major expenditure of resources! If the prima facie plausibility of the mediocrity principle rests, as I claim, on a tacit appeal to the principle of indifference, then ETI advocates might as well skip the mediocrity trimmings and just make a direct appeal to indifference. If the argument from mediocrity were valid, it would be superfluous.

To summarize the last three sections: there are no arguments on the table that would underwrite any degree of optimism—or pessimism—about the ETI hypothesis. The same applies to the stronger SETI and CETI hypotheses. The conclusion isn't that we should repudiate these hypotheses. It's that we have no grounds for having any opinion at all. A few words of clarification on what this epistemic position is like. So long as we describe the epistemic options in terms of the dichotomous folk-psychological concept of belief, the distinction between disbelief and

agnosticism is clear enough. To disbelieve a hypothesis X is to believe that its negation not-X is true; to be agnostic about X is neither to believe X nor to believe not-X. When we talk in terms of probabilities, however, the agnostic option of suspending judgment may seem to be lost. What is the probability of X for an agnostic? To say that it's 1/2 is to appeal illegitimately to the principle of indifference. But there's also no warrant for making the probability of X either greater than or less than 1/2. To make sense of agnosticism, we have to make a conceptual space for *vague* probabilities.[6] Hypotheses can't always be associated with discrete probabilities. Sometimes the best we can do is to ascribe a probability *interval* to the state of affairs under consideration. For example, we may know nothing more than that X is likelier to be true than false. This epistemic state of affairs would be represented by ascribing the vague probability (1/2, 1] to X. (A parenthesis indicates that the end-point is excluded from the interval, while a square bracket indicates that the end-point is included in the interval.) To say that we have no idea whether X is true or false is to associate X with the entire probability interval [0, 1]. My assessment of the state of the three extraterrestrial hypotheses can be represented by saying that there's no basis for associating anything less with them than the full probability interval [0, 1].

5. Decisions Under Conditions of Pervasive Uncertainty

In this section and the next, I'll take a closer look at the metaquestion whether ETI is a project worthy of intensive pursuit. In the past three sections, we've tried to ascertain the probability that the pursuit will meet with success. But obviously the answer to the metaquestion depends on both the likelihood of success and on how interesting or important a successful outcome would be. A high probability of success wouldn't make a project pursuitworthy if the cost were enormously high and the payoff relatively puny. For example, we'd have a high probability of succeeding if we were determined to record the license plates of all the cars that pass by a randomly chosen crossroad in Little Rock, Arkansas for a period of twenty years. But that doesn't mean it's a good idea to do it. Conversely, a low probability of success—even a very low probability of success—may be compatible with pursuitworthiness if the potential payoff is enormous.

How do probabilities and values combine to determine a rational decision? This isn't the place to entertain novel proposals for measures of program evaluation. Throughout this article, my aim is to ascertain where ETI stands in relation to the currently received canons of scientific rationality. I will approach the metaquestion from the standpoint of traditional decision theory. Or rather, I'll do so as much as I can. It will turn out, however, that the pervasive absence of information with which the decision has to be made forces us to introduce some conceptual exten-

[6]An excellent account of vague probabilities is given by Bas van Fraassen (1989).

sions to the standard decision-theoretical framework erected by von Neumann and Morgenstern. However, these extensions are natural, inevitable, and I think uncontroversial.

In decision theory, it's conceived that the decision-maker is confronted with a number of mutually exclusive courses of *action*. The task is to provide rational grounds for selecting one of them. Each possible action is conceived to have any number of possible *outcomes*, and each outcome has assigned to it a *probability* and a *value*. Let $x_1, x_2, \ldots, x_n$ be the n possible outcomes of action X. Then $p(x_i)$ is the probability that x_i will occur as a consequence of choosing to do X. The value $v(x_i)$ of x_i is a measure of the relative preference for this outcome as compared to others. There are well known decision-theoretical techniques for assigning numbers to values up to an *interval* scale, which is to say that the zero-point and unit value are arbitrarily selected, whereupon all the other values are fixed. Decision theory stipulates that the best course of action is the one with the highest *expected value* (EV), this quantity being the sum of the products of the probabilities and values of all that action's outcomes:

$$EV(X)=p(x_1)v(x_1)+p(x_2)v(x_2)+\cdots+p(x_n)v(x_n).$$

Why should we maximize EV? Because a persistent policy of doing so results in a greater realization of value than any other policy in the long run.

In most cases of real-life decision-making, one of the available courses of action is *inaction*—doing nothing at all. Inaction can be treated as just another action to which an EV is assigned on the basis of its possible outcomes and their values and probabilities. The outcomes of inaction may be substantial. If a gunman threatens to shoot you unless you play either game X or game Y, the EV of inaction will undoubtedly be less than EV(X) and EV(Y), and you would do well to choose a game and play it even if they're both boring. If no particular consequence is expected to follow from inaction, then inaction is a choice with a single outcome—the continuation of the status quo. The continuation of the status quo will of course have a value of its own, which may be either higher or lower than the values of the outcomes associated with other available courses of action. For example, the status quo may include the fact that I have a toothache, in which case inaction could have a lower EV than the option of going to the dentist.

Choices with a single outcome come up in other ways too. For instance, you may be given a choice between playing a game of chance and receiving a certain prize outright. If you opt for the outright prize, then it's yours for sure. In such a case, the action of accepting the prize has a single outcome x, whose probability p(x) of occurring if that action is taken is 1. Thus the expected value of that choice is:

$$EV=p(x)v(x)=1v(x)=v(x).$$

That is to say, the expected value of a single-outcome choice is identical to the

value of that outcome. If inaction has no other possible outcome except the continuation of the status quo, then the expected value of inaction is the value of the status quo. Since the zero point of value is arbitrary, it's often convenient to assign the status quo a value of zero.

The analysis of the ETI hypothesis in the previous sections led to the conclusion that its probability is radically vague. How do vague probabilities fit into the decision-theoretical picture? Consider a game of chance X where the probabilities of the two possible outcomes—call them w and l for 'win' and 'lose'—are totally vague. This means that p(w) may be any number between 0 and 1. Whatever p(w) may be, it has to be the case that $p(w) + p(l) = 1$. So while both p(w) and p(l) can range from 0 to 1, they can't both be very low or very high simultaneously. Now suppose that the outcome w has a greater value than the outcome l (that's why we call it 'w'). Then the indeterminacy in probabilities entails a corresponding indeterminacy in the expected value of the game X. The highest value that EV(X) can have is obtained when $p(w) = 1$ and $p(l) = 0$. In this case,

$$\max EV(X) = p(w)v(w) + p(l)v(l) = 1v(w) + 0v(l) = v(w).$$

The lowest value that EV(X) can have is obtained when $p(w) = 0$ and $p(l) = 1$, whereupon

$$\min EV(X) = p(w)v(w) + p(l)v(l) = 0v(w) + 1v(l) = v(l).$$

Thus EV(X) can be anything between the value of winning and the value of losing. If we can't be more precise about probabilities, then we also can't be more precise about expected values.

Nevertheless, this information may be enough to make a rational decision. If you play game X, the worst-case scenario is that its EV is equal to v(l), the value of losing. If your other options all have EVs *less* than v(l), then X is the indicated choice, despite the radical indeterminacy in its expected value. For example, if somebody offers to play a (painless) game with you in which you receive $2 if you win and you receive $1 if you lose, you would do well to accept the offer even if you know nothing at all about the odds of winning. This configuration of circumstances actually comes up in the evaluation of the pursuitworthiness of ETI. We'll see that some ETI advocates maintain that the pursuit of ETI will pay off handsomely even if it fails—for failure, like success, would tell us some very important things about ourselves and the universe we live in. Call this the *can't-lose* hypothesis. If the can't-lose hypothesis is true, then of course we can answer the metaquestion in the affirmative even though the probabilities of ETI success or failure are entirely indeterminate.

It's also possible to make a decision in the face of probabilistic indeterminacy if the values of both the winning and the losing outcome are less than the expected value of another available alternative. In this case, the decision is that we should *not* undertake the indicated course, for we can do better no matter what the outcome

turns out to be. In all other cases, however, probabilistic indeterminacy robs decision theory of the basis for making a recommendation. Suppose that the value of winning is higher than the expected value of the best alternative action, but that the value of losing is lower than the expected value of the worst available alternative. Then there's no way of knowing whether we should play the game. The decision has to be made non-rationally.

We're almost ready to assess the values of success and failure at ETI. There's just one more preliminary issue: an argument by Regis (1985) to the effect that failure is not one of the possible outcomes of engaging in ETI research. Regis maintains that we can never arrive at an epistemic state from which it's appropriate to say that ETI has failed. His point isn't merely that the probability of failure, like the probability of success, is radically indeterminate. It's that the failure of ETI can never be established, no matter what additional information may come our way in the future. In this respect, there's an asymmetry between ETI success and ETI failure. However likely or unlikely they may be, there are scenarios that would persuade us of the truth of the ETI hypothesis. Suppose, for example, that we track a spaceship travelling from beyond the orbit of Pluto until it lands on earth, and that five-legged beings come out of the ship and begin to converse with us in English, explaining that they come from a planet circling the star of Sirius and that they've learned English by analyzing our old radio and TV broadcasts. Now there may be good reasons for believing that this scenario will never be actualized. But if it *were* to be actualized, then we would know that ETI (as well as SETI and CETI) is true. Regis claims that we can't in a similar manner envision an ascertainable state of affairs that would definitively indicate that ETI is false:

> The problem is that even a major search's failure to detect an interstellar message would not by itself disprove the existence of extraterrestrials. For no matter how extensive and complete a search may be, the possibility would always remain that it was not complete enough, that maybe we had not searched the right places at the right times, or on the right frequencies, or with the right reception media. There are countless things that we might have missed, or even have failed to look for: gravity waves, tachyon messages, neutrino modulations, and so on and so forth. Since there is no way of knowing that any given search—or series of them—has exhausted all the possibilities, there is always the residual chance that undetected aliens are still out there, that whole civilizations have escaped our notice . . . [I]f we make a search and end up only with absence of evidence for ETI, then the only thing we would be entitled to conclude from this is that we have no evidence of their absence. In other words, we are no better off—and no worse off either—than we were when we started: we still don't know if they're out there. (Regis, 1985, p. 233)

If Regis is right, then the two outcomes that jointly determine the expected value of ETI aren't success and failure, but success and no resolution at all. Contrary to the suggestions of some ETI advocates, non-success wouldn't teach us anything about the rarity or preciousness of intelligent life. It wouldn't teach us anything at all. The alternative to success is in fact a continuation of the epistemic status quo.

But the continuation of the status quo is free. Therefore, unless there is some prospect for success, ETI isn't worth a dime.

The confirmational problem that Regis points to affects all non-existence claims in the same way. A hypothesis to the effect that a certain type of entity or process *exists* might come to be firmly established on the basis of a single serendipitous observation. This isn't the case with the hypothesis that a certain type of object *doesn't* exist—for no matter how many places and times the object is observed to be absent, it remains possible that it exists somewhere else or at some other time. Nevertheless, Regis' conclusion is too severe. Despite the difficulty of verifying negative existential claims, many propositions of this type make their way into science. Contemporary science affirms the non-existence of the luminiferous aether, phlogiston, unicorns, Abominable Snowmen, perpetual motion machines, and so on. Some of these beliefs are based on ordinary inductive reasoning. While it's true that we haven't examined every place where Abominable Snowmen might lurk, the fact that they haven't been located in so many of the places where they might have been gives us inductive grounds for supposing that they aren't in any of the unexamined locales either. To be sure, we might be wrong: inductive inference always brings the possibility of error in its train. But inferences to universal *absences* are not riskier than inferences to universal *presences*. The move from the absence of Abominable Snowmen or ETIs in the places we've looked to their absence in places we haven't looked is on the same epistemic footing as the move from the blackness of observed ravens to the blackness of unobserved ravens.

More importantly, Regis doesn't take into account the possibility of a *theoretical* resolution to the ETI question. Let's grant that our failure to find ETIs can never warrant the belief that they don't exist. Even so, we might arrive at the conclusion that ETIs don't exist on the basis of a theoretical argument. Theoretical reasoning is the basis for many of the non-existence claims made by scientists. Physicists' conviction that there can never be a perpetual motion machine isn't primarily based on an inductive generalization from the fact that nobody has ever succeeded in building one—it's based on the laws of thermodynamics. By the same token we may come to have theoretical reasons for supposing that ETIs don't exist either, or that we can never find them even if they do exist, or that we can never communicate with them even if we find them. Whether by brute induction or by subtle theoretical reasoning, it's conceivable that we might rationally come to the conclusion that one or the other of the extraterrestrial projects is an irremediable failure.

6. The Metaquestion

What are the sources of value in ETI? I'll deal with this question in two stages. First I'll erect a rough-and-ready taxonomy of relevant types of value; then I'll assess how much of each type is actually possessed by ETI.

To begin with, there's the *epistemic* value of ETI—the value of what ETI success

or failure can teach us about the world. The main contrast to epistemic value is the *pragmatic* value of gifts other than knowledge that may be bestowed on us by well-intentioned aliens: new technologies, new art forms, new jokes, and so on. The gift of new sciences is intermediate between the two types of value. It's like epistemic value in that it amounts to the acquisition of new knowledge. But it's unlike other sources of epistemic value in two respects, one of which is important to the assessment of the metaquestion. By definition, the prototypical epistemic value is the value of the knowledge that *we* derive by exercizing our rational faculties on the datum that ETI (or SETI, or CETI) succeeds or fails. Receiving the gift of a new science is obviously not the same thing. But this is the unimportant difference. The reason why the gift of new science should be distinguished from prototypical epistemic values is that the former is not going to be forthcoming unless we succeed at *CETI* as well as SETI and ETI: we can't be given new sciences unless we're able to communicate with the would-be donors. In this regard, the gift of new sciences is akin to the gift of new art forms and jokes. In contrast, prototypical epistemic values may already be available with mere SETI success, or even with bare-boned ETI success: the bare fact that ETIs exist may already have major implications for our view of the world. For the purpose of assessing the metaquestion, the gift of new sciences will be grouped together with the gift of new jokes as a source of pragmatic value.

Another important distinction is drawn by Regis (1985). There's a difference between the epistemic consequences that can rationally be drawn from ETI success or failure, and the cognitive consequences that will in fact ensue, regardless of whether they're rationally warranted. For example, it has sometimes been suggested that ETI failure would produce the positive benefit of causing us to regard intelligent life as more precious than we currently do (Sagan, 1983). It's not at all clear that this is a rational inference. It seems to be based on the principle that the value of an intelligent life is a function of the total number of available intelligences. But it's a consequence of *this* principle that the murder of an individual human being becomes less reprehensible as the total population of the world increases. If we're not prepared to accept this consequence, then why should we view the non-existence of ETIs as relevant to the value of human life? This issue of moral theory aside, however, we might very well all agree that anything which increases the preciousness of human life is to be counted as advantageous. So if the non-existence of ETIs has that effect, then it is to be counted as a benefit that accrues to ETI failure, regardless of whether the effect is rationally justifiable. The distinction between this type of value and pragmatic value is that the latter involve the acquisition of *capacities* that we may or may not choose to exploit. The fact that we're given new sciences or new technologies or new jokes doesn't automatically entail that we have to accept, utilize or tell them, respectively. To the extent that the gift is only pragmatic, its employment or non-employment is up to us. In the case of ETI failure leading to an increase in the estimated preciousness of intelligent life,

however, the former stand in relation to the latter as cause to effect. Beneficial or harmful causal effects of ETI success or failure will be called *causal* values. In the example of ETI failure leading to an increase in the perceived preciousness of human life, the causal effect is cognitive: ETI failure produces a beneficial change in our mentality. The causal effects of the various ETI and SETI outcomes may also be non-cognitive. For example, SETI success may lead the contacted extraterrestrials to fly to earth and enslave us. If this is a real possibility, it would have to be considered in our overall assessment of the advisability of funding SETI projects. In my terminology, it would be a non-cognitive source of causal value.

Finally, there's the intrinsic interest in gratifying our curiosity about ETI, independently of its bearing on any broader issues or of its causal or pragmatic consequences. As Regis says:

> . . . SETI searches do not have to be justified on the grounds of legendary benefit . . . in order for it to be worth our while to continue searches . . . For there will be one undeniable benefit if one day we have a positive result: it will satisfy our curiosity to know what else is out there. (p. 243)

I call this the *gossip* value of ETI success or failure. Gossip value is a special case of epistemic value. It's the minimal amount of epistemic value that the discovery of any new fact always possesses. Some new facts possess enormous epistemic value by virtue of leading us to have new insights into the workings of the universe. They may, for example, crucially confirm or disconfirm a fundamental theory. Other new facts may not have much of a bearing on our other opinions. But every new fact possesses at least the epistemic value of causing us to learn that it itself is true. This irreducible minimum of epistemic value is that fact's gossip value. Different facts have different amounts of gossip value by virtue of being intrinsically more or less interesting. Everyone will undoubtedly agree that the gossip value of the putative fact that there are (or aren't) other intelligent beings in the universe is greater than the gossip value of the fact that the per capita consumption of cheese in the US was (or wasn't) 24 pounds in 1987.

Now let's consider cases: what are the values of success and failure at ETI, SETI and CETI? I've divided the topic into three parts: (1) the pragmatic and causal values of failure, (2) the pragmatic and causal values of success, and (3) the epistemic and gossip values of both failure and success. First what are the pragmatic and causal values of ETI failure? Since pragmatic values attach to capacities, there's no such thing as pragmatic harm. We might be indifferent to the acquisition of a capacity, but it can never diminish our fortune. If extraterrestrial humor causes us to cringe, we simply needn't indulge in it. So when we canvass the sources of value in the several extraterrestrial projects, we need only consider pragmatic values on the benefit side of the ledger. But it's clear that failure at any of the extraterrestrial projects isn't going to result in any increased capacity. Thus pragmatic values are irrelevant to ETI, SETI or CETI failure.

What about causal benefits? Some candidates have been floated. According to

Sagan, ETI failure 'would tend to calibrate something of the rarity and preciousness of the human species' (Sagan, 1983, p. 36); it 'would underscore, as nothing else in human history has, the individual worth of every human being' (Sagan, 1980, p. 314); it would have a 'sobering influence on the quarrelsome nation-states' (Sagan, 1982, p. 33). You get the idea. Presumably, this benefit would require full ETI failure. The more moderate failures of SETI or CETI wouldn't have the same sobering effect, since they wouldn't show that we're alone in the universe.

Regis has two counterarguments against Sagan's claim, one of which strikes me as persuasive. Both arguments make use of the historical fact that in pre-Copernican times European society was firmly convinced that human beings were the sole intelligent species in the universe. Regis notes that Sagan condemns this pre-Copernican view as a 'deadly chauvinism' (Sagan, 1983, p. 30). But how can the belief in our uniqueness be a good thing now if it was a bad thing then?

> [I]f it was wrong for the pre-Copernicans to have such distastefully ethnocentric, chauvinistic, and pridefully self-important notions about themselves . . . then how is it somehow beneficial to return to those same notions once the possibility of extraterrestrials is ruled out? This would seem to be having our cake and eating it too. (Regis, 1985, p. 234)

Regis has a point. But there are charitable ways to read Sagan on this issue. Sagan could consistently maintain that the knowledge that we're alone in the universe would have a salutary effect, but that it's a deadly *epistemic* chauvinism to be certain that we *are* alone in the universe. On this account, Sagan condemns the pre-Copernicans for accepting an improbable thesis, but he avers that if the improbable thesis were to turn out to be true, it would have good effects.

This reading of Sagan's thesis saves it from Regis' first charge of inconsistency. But it lays the thesis wide open to Regis' second charge of being counter-indicated by the facts in our possession. For the pre-Copernicans provide us with evidence of what the causal consequences actually are of believing that we are unique in the universe. And the evidence can hardly be said to support Sagan's conjecture:

> [I]f we examine the claim that failure to find aliens 'would have a sobering influence on the quarrelsome nation-states', we would have to wonder why the pre-Copernican conception of man's specialness failed to have this effect in years gone by. If, as seems to be true, the nations of those times were not more law-abiding and peaceful than modern states, then what reason do we have for imagining that a return to the old belief in humanity's specialness would have any such sobering influences today? This seems to be just wishful thinking. (Regis, 1985, p. 234)

A Saganite might protest that conditions are so different nowadays that there's little reason to believe that our reactions to the acceptance of human specialness would be the same as those of the pre-Copernicans. That may very well be. The evidence is far from conclusive. Nevertheless, it's true that what evidence we have, imperfect as it is, speaks against Sagan's hypothesis. But then a rational person should guess that Sagan's hypothesis is false. At best, it's a conjecture that strikes us as not entirely implausible. Now intuitive non-implausibility is *something*—all

other things being equal, we should probably prefer a hypothesis that isn't intuitively implausible to one that *is* intuitively implausible. But in this case, there are equally plausible hypotheses that lead to contrary conclusions. For example, we might not implausibly guess that the discovery that we're alone in the universe would precipitate humanity into a state of neurotic depression, loneliness and anomie. The historical evidence of pre-Copernican behavior doesn't support such a hypothesis, but that still leaves it in no worse an epistemic state than Sagan's more optimistic view.

But what is happening to our canvass of values? Decision theory tells us that the rational advisability of a course of action depends on both the probabilities and the values of the outcomes. We had dealt with probabilities in Section 2, Section 3 and Section 4, and now we were supposed to turn to considerations of value. But we find ourselves talking about probabilities once again—such as the likelihood that humanity's reaction to ETI failure would be an increase in the preciousness of human life or a suicidal depression. What's happening is that the idealizations of classical decision theory are proving to be inadequate for the representation of real-life decisions. Decision theory presumes that the values of the several possible outcomes are fully determinate: if you succeed, you definitely get $5, and if you fail, you definitely lose $1. But here the values of the outcomes—ETI success or ETI failure—are themselves uncertain, on account of our uncertainty concerning the *consequences* of success or failure. Since there's some question whether ETI failure will result in an increase in the preciousness of human life or a suicidal depression, we can't simply assign a value to this outcome. To deal with this more complex situation, we have to extend the conceptual apparatus of decision theory. We have to regard the outcomes themselves as having *expected* values, the calculation of which depends on the probabilities and values of their several possible consequences.

So what's the expected value of ETI failure? Let's look at a single possible consequence of failure: that there will be an increase in the preciousness of human life. Call this consequence C_1. The contribution of C_1 to the value of ETI failure is equal to the sum of two factors: (1) the probability $p(C_1)$ that C_1 will occur if ETI fails multiplied by the value $v(C_1)$ of that consequence, and (2) the probability that C_1 will *not* occur times the value of its *non*-occurrence. A similar pair of factors determines the contribution of the consequence C_2 that there will be a planet-wide depression. We may suppose that there are other states of affairs C_3, C_4, . . ., C_n, which are also possible consequences of ETI failure. In all these cases, it's appropriate to consider the non-occurrence of the consequence as equivalent to the continuation of the status quo, and to assign it a value of zero. If we do this, and if we assume that the several outcomes are probabilistically independent of one another, then the expected value of ETI failure can be represented as the sum of the products of the probabilities and values of each possible consequence:

$$EV = p(C_1)v(C_1) + p(C_2)v(C_2) + \cdots + p(C_n)v(C_n).$$

Of course, to assign discrete values to consequences like 'an increase in preciousness' and 'depression' is still to idealize the real situation—for the values of both these consequences obviously depend on the *degree* of preciousness or depression that's at issue. It's one thing if ETI failure produces universal and permanent disarmament; it's quite another thing if it merely affects the sentiments expressed in print by scientists and literati. If it were worth our while, this situation could be rigorously represented in the decision-theoretical scheme by drawing continuous functions that map degrees of preciousness to probabilities in such a way that the integral of the function over all degrees of preciousness is equal to 1. But given the rough-and-ready nature of our assessment of the relevant probabilities, to follow such a procedure would be to swat a conceptual fly with a conceptual sledgehammer. For our purpose, the several consequences of ETI failure can be represented as dichotomous: either the preciousness of human life increases, or it doesn't, and so on. Furthermore, we needn't include *all* the possible consequences of ETI failure in our considerations. It's enough if we take into account those that have a non-negligible probability of occurring *and* that have a substantial positive or negative value (taking zero as the value of the status quo). For example, it might become clear that ETI failure will quite likely produce an increase in the frequency with which words starting with the letter 'f' are used by speakers of English. But who cares?

So what are the probabilities and values of C_1, C_2, and any other relevant consequences of ETI failure? As for their probabilities, we don't face the same sort of blanket indeterminacy as we did when we tried to assess the likelihood of ETI failure itself. For what we're trying to estimate the probability of isn't a hypothesis about utterly unknown extraterrestrial conditions—it's a hypothesis about *our own* reactions to certain beliefs which have in fact previously been held by some societies in some eras. I refer once again to the pre-Copernican cosmology of Europe. The evidence is that this society's belief in the uniqueness of human intelligence didn't produce an increase in the perceived preciousness of human life, or in Weltschmerz, or in any other social good or evil. Once again, the evidence is far from conclusive. But it's the only basis we have for formulating an opinion about the $p(C_i)$s. The best hypothesis is thus that none of the C_i will happen. But if they won't happen, then we don't have to agonize over what values we should assign to them: the expected value of ETI failure is going to be zero regardless of what the values of the C_i are.

What about SETI and CETI failure? Well, ETI failure *entails* SETI failure—we won't succeed in contacting extraterrestrials if none of them exists. Thus any source of value that accrues to SETI failure must also accrue to ETI failure. But we've just seen that *no* pragmatic or causal values accrue to ETI failure. Therefore no pragmatic or causal values accrue to SETI failure either. Similarly, SETI failure entails CETI failure. A reiteration of the same argument shows that CETI failure

has no pragmatic or causal value. The grand conclusion is that (epistemic and gossip value aside) Sagan's 'can't-lose' hypothesis must be rejected. The only hope for establishing the worthiness of any of the extraterrestrial projects is by an appeal to the benefits of their success. If we expend any resources on the extraterrestrial projects and fail, we will have lost our investment.

Let's turn now to the value of success. I'll start with the minimal case of ETI success without accompanying SETI or CETI success—i.e., just finding out that there are ETIs, without actually coming into contact with them. In this minimal case, there are no pragmatic benefits to be obtained. However, Sagan argues that ETI success would have the causal benefit of providing 'a powerful integrating influence on the nations of the planet earth. [It] would make the differences that divide us down here on earth increasingly more trivial' (Sagan, 1982, p. 32). This putative benefit seems to be pretty much the same as the putative benefit of ETI *failure*, which Sagan surmised would have a 'sobering influence on the quarrelsome nation-states' (Sagan, 1982, p. 33). Regis argues that Sagan can't have it both ways: if the discovery that we're alone in the universe diminishes the likelihood of war, then shouldn't the discovery that we're *not* alone *increase* the likelihood of war?

> After all, if we should blow ourselves to pieces here on earth, there is another whole civilization out there waiting to take up where we left off. Made complacent by the knowledge that other intelligences exist, we can play fast and loose with the fate of the earth. (Regis, 1985, p. 237)

Sagan's conjectures and Regis' critique can be understood either epistemically or causally: it may be claimed that ETI success (or failure) would provide pacifists with additional support for their position, or that it would *cause* us willy-nilly to become more pacifistic. Regis' objection is telling against the epistemic thesis, although Regis doesn't provide us with all the steps of the argument. He seems to suggest that Sagan's position is logically inconsistent—that you can't, on logical grounds alone, assert that Q is a consequence of both P and not-P. In fact, you *can* assert that Q is a consequence of both P and not-P. But if you do, then the truth-value of P is irrelevant to your acceptance of Q—you should accept Q even if you're totally in the dark about P. If both the discovery that we're alone and the discovery that we're not alone rationally entail pacifism, then we should be pacifists right now, since one or the other of these two alternatives is bound to be correct. So far as the pacifism issue is concerned, there's nothing to be gained by waiting to find out whether we're alone. ETI makes no difference to the case for pacifism.

What if we interpret Sagan as making the causal claim that both ETI success and ETI failure would incline us, whether rationally or irrationally, toward pacifism? In this case, Regis' objection is inconclusive. Pointing out that Sagan assigns the same effect to ETI success as he does to ETI failure may serve to undermine the intuitive plausibility of Sagan's pair of conjectures. But it's at least logically possible that

both causal conjectures are true—that both the discovery that ETIs exist and the discovery they don't exist might have a pacifying influence on the quarrelsome nation-states. But the causal conjecture about ETI success is liable to the same pair of objections as the causal conjecture about ETI failure. First, it's no more plausible than the opposite guess that the discovery of ETIs will, as Regis suggests, make us more complacent about blowing ourselves to pieces. Secondly, what evidence we have doesn't support either guess:

> We have prior experience to guide us here as well, the finding of new worlds right here on earth. The discovery of the Americas, for example, did not have anything like the effect upon Europeans that SETI advocates insist that discovery of ETI will have upon us. It did not make differences among Europeans more trivial, it did not serve as an integrating influence among them, it did not make them more tolerant and peace-loving. (Regis, 1985, p. 237)

So, epistemic and gossip value aside, our best guess has to be that ETI success without accompanying SETI or CETI success has no value. Since ETI failure is also valueless, the conclusion is that the minimal project of ETI sans SETI or CETI is worthless.

On the face of it, it's CETI success that might secure for us a pragmatic bonanza of new sciences and new arts. However, Regis argues that the pragmatic value of CETI success is vanishingly small. For suppose that we receive a message which conveys an account of a new ethical or political system. How shall we assess its worth?

> The only way to assess alien moral or political systems is to compare them to the systems of earth, and to project the consequences of adopting their doctrines here. Such an analysis would in the end lead us to embrace those aspects of their ethics that we thought would further human well-being, and to reject those that would not. But in doing this we would be using our own pre-contact standards. Call it provincial, reactionary, chauvinistic, or what you will, we'd be evaluating other theories by reference to our own. But to do this, to use our own accepted moral or political standards as criteria for the evaluation of other, alien ideologies, is hardly to get outside the earthbound perspective from which alien communications are supposed to liberate us. (Regis, 1985, p. 240)

And so it goes, presumably, with other potential pragmatic benefits: we'll never learn anything radically new from ETIs, because we'll have to evaluate everything they tell us by our own terrestrial lights. The tip-off that there's something wrong with this argument is that its import is far too strong. Substitute two arbitrary parties for earthlings and aliens, and it becomes an argument for the proposition that nobody can ever learn anything new from anybody else. Here's a quick (but far from complete) account of what's wrong with Regis' argument. New cultural forms, whether moral, political, aesthetic or scientific, arise in two stages: first somebody comes up with a new idea, and then the new idea is evaluated as worthy of implementation. All of Regis' remarks address the second evaluative stage. Let's suppose that he's right in claiming that our evaluative criteria are unalterable. It

doesn't follow from this that extraterrestrials can't teach us anything new, for they may still present us with new ideas that have both of the following properties. First, by virtue of their different biology, different planetary environment, and different history, certain ideas might occur to ETIs that would never spontaneously occur to any of us humans. Secondly, when these ideas, which we ourselves would never abduce, are presented to us by an external agency, we might find that some of them rank very high *according to our own system of evaluation.* Nevertheless, we might never have thought of them ourselves. Assuming that CETI succeeds, there's no reason to believe that we couldn't profit from a Plutonian's distinctive take on the universe in this way. Of course, we can't be sure that ETIs will be both willing and able to provide us with pragmatic benefits. But since there can be no pragmatic harm (pragmatic values being capacities), we have everything to gain and nothing to lose by CETI success—pragmatically speaking.

On the other hand, CETI success could have catastrophic *causal* consequences that counterbalance the possible pragmatic benefits. If the ETIs make contact with us, they might come to earth in order to enslave us, or cook us for dinner. Considering both pragmatic and causal values together, the possible consequences of CETI success range from the realization of our fondest fantasies to the actualization of our worst nightmares. The same is true of SETI success. ETIs could cook us for dinner whether or not we had succeeded in communicating with them—and while they can't convey new sciences to us without CETI success, they could still be the causes of wonderful effects. For example, they could innoculate us against all diseases even if they couldn't communicate with us, just as we innoculate dogs and cats.

But how likely are these fantasies and nightmares? When we evaluated the causal effects of ETI success and failure, we found reason to believe that the fantasy of universal disarmament and the nightmare of universal Weltschmerz were very low-probability consequences. In these cases, we had something to go on because the hypothesized consequences were hypotheses about how *we* would behave under conditions that were not entirely unlike conditions that we've faced in the past. The hypothesized consequences of SETI and CETI success, however, are hypotheses about how *aliens* would behave. Clearly, we have no more basis for an informed opinion about these matters than we have on the probability of ETI, SETI or CETI success itself. The result is that the value of SETI or CETI success is itself radically indeterminate. It could range from the greatest imaginable boon to the greatest imaginable curse—from immortality to being cooked and eaten.

So what's the expected value of SETI or CETI? It's given by the formula

$$EV = p(s)v(s) + p(f)v(f),$$

where p(s) and p(f) are the probabilities of success and failure, and v(s) and v(f) are the values of success and failure. Earlier in this section, I argued that the value of SETI or CETI failure is zero. Thus the expected value of SETI or CETI as a

whole is equal to the product of the probability of success and the value of success. But both of these quantities are radically indeterminate. In Section 4, we saw that p(s) is vague over the entire interval from 0 to 1. In the present section, we saw that v(s) can range from the highest good that we can conceive to the worst evil that we can conceive. The upshot is that the expected value of SETI or CETI itself is indeterminate between the highest good and the worst evil. Everything that we know is compatible with our receiving the gift of immortality with a probability of 1 if we engage in SETI, or with our being cooked and eaten with a probability of 1, or with anything in between.

Thus rationality provides us with no guidance for how to answer the metaquestion about SETI or CETI. Spending a fortune on trying to make contact with extraterrestrials has as much rational warrant (namely, none) as spending a fortune on trying to *avoid* contact with extraterrestrials. Everybody's predilection is as good as everybody else's. To be sure, we've only considered pragmatic and causal values so far. But this null result isn't altered when we bring epistemic and gossip values into the picture. In the first place, it was argued in Section 2 that the several extraterrestrial projects have not been shown to have a bearing on any theoretical issues. In this regard, the fact that there are (or aren't) any ETIs is very different from the fact that the perihelion of the planet Mercury revolves about the sun by an amount of 43 seconds of an arc every hundred years. The latter gives us an excellent reason to accept Einstein's paradigm-breaking views of space, time and gravitation. The former doesn't change much of anything in our cognitive world. But this is to say that the extraterrestrial projects have no epistemic value to speak of. If tomorrow we received a radio message from space that was clearly of artificial origin, there would be huge headlines in the newspapers for a few days or weeks or months. But when the dust had cleared, none of our scientific ideas about the universe would have been changed. As things stand at present, the search for extraterrestrial intelligence is scientifically *boring*.

This conclusion is sometimes denied by biologists. For example, Mayr (1985) seems to suggest that evolutionary theory leads to the expectation that there are no ETIs. If this is right, then the discovery of ETIs would be an unexplained mystery calling for revisions in our received views about evolution. But evolutionary theory has nothing to do with Mayr's anti-ETI conclusion. Mayr's argument is purely inductive: of the billions of species that have existed on earth, only one has possessed sufficient intelligence to engage in SETI (see Section 2). It isn't evolutionary theory that tells us that intelligent species are improbable. The force of Mayr's argument would be unchanged if, instead of believing in Darwinian evolution, we thought that all the species were created by special acts of God. The inductive argument would then be that ETI is improbable because the available evidence tells us that God hardly ever chooses to produce intelligent species. Establishing that there are (or aren't) ETIs neither confirms nor disconfirms the hypothesis of Darwinian evolution. It merely adds a detail to our natural history of the universe.

So bringing epistemic values into the picture doesn't alter the answer to the metaquestion. It only adds zeros to the zero value of minimal ETI success and to the enormous range of positive and negative pragmatic/causal values associated with SETI and CETI success. That leaves us with the gossip value of the outcomes of the several extraterrestrial projects. There's no doubt that the facts about ETI have a lot of gossip value—almost as much, perhaps, as the facts about the O. J. Simpson murder trial. This is reason enough for engaging in a minimal ETI project—a project of thinking about whether ETIs exist *without* trying to make contact with them. When it comes to SETI or CETI, however, the addition of a middling positive value to the maximally wide range of pragmatic and causal values once again leaves the indeterminacy unchanged. We couldn't decide whether SETI was pursuitworthy because SETI success might either make us immortal or get us eaten. Now we realize that SETI success might either (1) make us immortal and gratify our curiosity, or (2) get us eaten and gratify our curiosity. The value is not less indeterminate than before. The interval of values has merely shifted a notch toward the positive. So long as the positive gossip value of SETI success has a smaller absolute magnitude than the negative value of the destruction of all life on earth, the answer to the metaquestion will still be indeterminate.

Finally, it's worth mentioning that the answer to the SETI metaquestion would remain indeterminate even if we were to get a fix on the probability of SETI success. To the extent that SETI is likely to fail, the range of expected values for SETI diminishes. But the metaquestion remains unresolved because, for all p(s) other than $p(s) = 0$, the range of associated values continues to straddle both sides of zero (i.e., the value of the status quo). As the probability of success diminishes, the stakes get smaller, but the decision whether to pursue SETI doesn't become any easier to make. With $p(s) = 0.0001$, we merely take a minuscule risk of being exterminated by aliens if we pursue SETI. But we also obtain a minuscule chance of being granted immortality. So there's no rational resolution to the metaquestion all the way down to $p(s) = 0$. When p(s) is exactly zero, the metaquestion whether to pursue SETI disappears, for there's nothing left to pursue. We've already got our answer: SETI fails. So for all probabilities of success, the question whether to pursue SETI either has no answer, or it doesn't arise. The same can be said of CETI.

References

Baird, J. C. (1987) *The Inner Limits of Outer Space* (Hanover: University Press).

Barrow, J. D. and Tipler, F. J. (1988) *The Anthropic Cosmological Principle* (Oxford: Oxford University Press).

Beck, L. W. (1985) 'Extraterrestrial Intelligent Life', in E. Regis Jr. (ed.), *Extraterrestrials: Science and Alien Intelligence* (Cambridge: Cambridge University Press), pp. 3–18.

Chow, S. and Teicher, H. (1978) *Probability Theory: Interdependence, Interchangeability, Martingales* (New York: Springer-Verlag).

Clarke, A. C. (1973) *Report on Planet Three* (New York: Signet).

Crowe, M. J. (1986) *The Extraterrestrial Life Debate 1750-1900: The Idea of a Plurality of Worlds from Kant to Lowell* (Cambridge: Cambridge University Press).
Davidson, D. (1973–4) 'On the Very Idea of a Conceptual Scheme', *Proceedings and Addresses of the American Philosophical Association* **67**, 5–20.
Dick, S. J. (1982) *Plurality of Worlds: The Origins of the Extraterrestrial Life Debate from Democritus to Kant* (Cambridge: Cambridge University Press).
Dick, S. J. (1996) *The Biological Universe: The Twentieth-Century Extraterrestrial Life Debate and the Limits of Science* (Cambridge: Cambridge University Press).
Dole, S. H. (1964) *Habitable Planets for Man* (New York: Blaisdell).
Hart, M. H. (1982) 'Atmospheric Evolution, the Drake Equation, and DNA: Sparse Life in an Infinite Universe', in M. H. Hart and B. Zuckerman (eds), *Extraterrestrials — Where Are They?* (New York: Pergamon Press), pp. 154–165.
Jones, H. S. (1940) *Life on Other Worlds* (New York: Mentor).
Kukla, A. (1991) 'Criteria of Rationality and the Problem of Logical Sloth', *Philosophy of Science* **58,** 486–490.
Lucretius (1951) *On the Nature of Things* (London: Penguin Books).
Mash, R. (1993) 'Big Numbers and Induction in the Case for Extraterrestrial Intelligence', *Philosophy of Science* **60,** 204–222.
Mayr, E. (1985) 'The Probability of Extraterrestrial Intelligent Life', in E. Regis Jr. (ed.), *Extraterrestrials: Science and Alien Intelligence* (Cambridge: Cambridge University Press), pp. 23–30.
McDonough, T. R. (1987) *The Search for Extraterrestrial Intelligence: Listening for Life in the Cosmos* (New York: Wiley).
McMullin, E. (1980) 'Persons in the Universe', *Zygon* **15,** 69–89.
Paine, T. (1948) *The Age of Reason* (New York: Liberal Arts Press).
Regis, E. Jr. (1985) 'SETI Debunked', in E. Regis Jr. (ed.), *Extraterrestrials: Science and Alien Intelligence* (Cambridge: Cambridge University Press), pp. 231–244.
Rescher, N. (1985) 'Extraterrestrial Science', in E. Regis Jr. (ed.), *Extraterrestrials: Science and Alien Intelligence* (Cambridge: Cambridge University Press), pp. 83–116.
Rood, R. T. and Trefil, S. J. (1981) *Are We Alone? The Possibility of Extraterrestrial Civilizations* (New York: Scribners).
Sagan, C. (1980) *Cosmos* (New York: Random House).
Sagan, C. (1982) 'The Search for Who We Are', *Discover* **3,** 31–33.
Sagan, C. (1983) 'We Are Nothing Special', *Discover* **4,** 30–36.
Salmon, W. C. (1966) *The Foundations of Scientific Inference* (Pittsburgh: University of Pittsburgh Press).
Shklovskii, I. S. and Sagan, C. (1966) *Intelligent Life in the Universe* (San Francisco: Holden Day).
Tipler, F. J. (1985) 'Extraterrestrial Intelligent Beings do not Exist', in E. Regis Jr. (ed.), *Extraterrestrials: Science and Alien Intelligence* (Cambridge: Cambridge University Press), pp. 133–150.
van Fraassen, B. (1989) *Laws and Symmetries* (Oxford: Clarendon Press).
von Mises, R. (1957) *Probability, Statistics and Truth* (New York: Macmillan).